arXiv:1512.02690v1 [q-bio.MN] 8 Dec 2015

Conditions for duality between fluxes and concentrations in biochemical networks

Ronan M.T. Fleming[a,*], Nikos Vlassis[b], Ines Thiele[a], Michael A. Saunders[c]

[a]*Luxembourg Centre for Systems Biomedicine, University of Luxembourg, 7 avenue des Hauts-Fourneaux, Esch-sur-Alzette, Luxembourg.*
[b]*Adobe Research, 345 Park Ave, San Jose, CA, USA.*
[c]*Dept of Management Science and Engineering, Stanford University, Stanford, CA, USA.*

Abstract

Mathematical and computational modelling of biochemical networks is often done in terms of either the concentrations of molecular species or the fluxes of biochemical reactions. When is mathematical modelling from either perspective equivalent to the other? Mathematical duality translates concepts, theorems or mathematical structures into other concepts, theorems or structures, in a one-to-one manner. We present a novel stoichiometric condition that is necessary and sufficient for duality between unidirectional fluxes and concentrations. Our numerical experiments, with computational models derived from a range of genome-scale biochemical networks, suggest that this flux-concentration duality is a pervasive property of biochemical networks. We also provide a combinatorial characterisation that is sufficient to ensure flux-concentration duality. That is, for every two disjoint sets of molecular species, there is at least one reaction complex that involves species from only one of the two sets. When unidirectional fluxes and molecular species concentrations are dual vectors, this implies that the behaviour of the corresponding biochemical network can be described entirely in terms of either concentrations or unidirectional fluxes.

Keywords: biochemical network, flux, concentration, duality, kinetics

1. Introduction

Systems biochemistry seeks to understand biological function in terms of a network of chemical reactions. Systems biology is a broader field, encompassing systems biochemistry, where understanding is in terms of a network of interactions, some of which may not be immediately identifiable with a particular chemical or biochemical reaction. Mathematical and computational modelling of biochemical reaction network dynamics is a fundamental component of systems biochemistry. Any genome-scale model of a biochemical reaction network will give rise to a system of equations with a high-dimensional state variable, e.g., there are at least 1000 genes in *Pelagibacter ubique* (Giovannoni et al., 2005), the smallest free-living microorganism currently

*Corresponding author.
Email addresses: ronan.m.t.fleming@gmail.com (Ronan M.T. Fleming), nikos.vlassis@gmail.com (Nikos Vlassis), ines.thiele@gmail.com (Ines Thiele), saunders@stanford.edu (Michael A. Saunders)

known. In order to ensure that mathematical and computational modelling remains tractable at genome-scale, it is important to focus research effort on the development of robust algorithms with time complexity that scales well with the dimension of the state variable.

Given some assumptions as to the dynamics of a biochemical network, a mathematical model is defined in terms of a system of equations. Characterising the mathematical properties of such a system of equations can lead directly or indirectly to insightful biochemical conclusions. Directly, in the sense that the recognition of the mathematical property has direct biochemical implications, e.g., the correspondence between an extreme ray of the steady state (irreversible) flux cone and the minimal set of reactions that could operate at steady state (Schuster et al., 2000). Or indirectly, in the sense of an algorithm tailored to exploit a recognised property, which is subsequently implemented to derive biochemical conclusions from a computational model, e.g., robust flux balance analysis algorithms (Sun et al., 2013) applied to investigate codon usage in an integrated model of metabolism and macromolecular synthesis in *Escherichia coli* (Thiele et al., 2012).

Mathematical duality translates concepts, theorems or mathematical structures into other concepts, theorems or structures in a one-to-one manner. Sometimes, recognition of mathematical duality underlying a biochemical network modelling problem enables the dual problem to be more efficiently solved. An example of this is the problem of computing *minimal cut sets*, i.e., minimal sets of reactions whose deletion will block the operation of a specified objective in a steady state model of a biochemical network (Klamt and Gilles, 2004). Previously, computation of minimal cut sets required enumeration of the extreme rays of part of the steady state (irreversible) flux cone, which is computationally complex in memory and processing time (Haus et al., 2008). By recognising that minimal cut sets in a primal network are dual to extreme rays in a dual network (Ballerstein et al., 2012), one can compute select subsets of extreme rays for the dual network that correspond to minimal cut sets with the certain desired properties in the primal (i.e., original) biochemical network in question (von Kamp and Klamt, 2014). This fundamental work has many experimental biological applications, including metabolic engineering (Mahadevan et al., 2015).

Recognition of mathematical duality in a biochemical network modelling problem can have many theoretical biological applications, in advance of experimental biological applications. For example, in mathematical modelling of biochemical reaction networks, there has long been an interest in the relationship between models expressed in terms of molecular species concentrations and models expressed in terms of reaction fluxes. When concentrations or *net* fluxes are considered as independent variables, a duality between the corresponding Jacobian matrices has been demonstrated (Jamshidi and Palsson, 2009). In this case, the concentration and net flux Jacobian matrices can be used to estimate the dynamics of the same system, with respect to perturbations to concentrations or net fluxes about a given steady state. The primal (concentration) Jacobian and dual (net flux) Jacobian matrices are identical, except that one is the transpose of the other. Matrix transposition is a one-to-one mapping and the aforementioned duality is between the pair of Jacobians. This does not mean that the net flux and concentration vectors are dual variables in the same mathematical sense, and neither are the perturbations to concentrations or net fluxes. This is because the Jacobian duality (Jamshidi and Palsson, 2009), which exists for any stoichiometric matrix, does not enforce a one-to-one mapping between concentrations and net fluxes unless the stoichiometric matrix is invertible, which is never the case for a biochemical network (Heinrich et al., 1978).

Herein we ask and answer the question: what conditions are necessary and sufficient for duality between unidirectional fluxes and molecular species concentrations? We establish a necessary

linear algebraic condition on reaction stoichiometry in order for duality to hold. We also combinatorially characterise this stoichiometric condition in a manner amenable to interpretation for biochemical networks in general. In manually curated metabolic network reconstructions, across a wide range of species and biological processes, we confirm satisfaction of this stoichiometric condition for the major subset of molecular species within each reconstruction of a biochemical network. Furthermore, we demonstrate how linear algebra can be applied to test for satisfaction of this stoichiometric condition or to identify the molecular species involved in violation of this condition. We also demonstrate that violation of flux-concentration duality points to discrepancies between a reconstruction and the underlying biochemistry, thereby establishing a new stoichiometric quality control procedure to select a subset of a biochemical network reconstruction for use in computational modelling of steady states.

First, we establish a linear algebraic condition and a combinatorial condition for duality between unidirectional fluxes and concentrations. Subsequently, we introduce a procedure to convert a reconstruction into a computational model in a quality-controlled manner. We then apply this procedure to a range of genome-scale metabolic network reconstructions and test for the linear algebraic condition for flux-concentration duality before and after conversion into a model. We conclude with a broad discussion, with examples illustrating how a recognition of flux-concentration duality could help address questions of biological relevance and improve our understanding of biological phenomena.

2. Results

2.1. Stoichiometry and reaction kinetics

We consider a biochemical network with m molecular species and n (net) reactions. Without loss of generality with respect to genome-scale biochemical networks, we assume $m \leq n$. We assume that each reaction is *reversible* (Lewis, 1925) and can be represented by a *unidirectional reaction* pair. With respect to the forward direction, in a *forward stoichiometric matrix* $F \in \mathbb{Z}^{m\times n}$, let F_{ij} be the *stoichiometry* of molecule i participating as a substrate or catalyst in *forward unidirectional reaction* j. Likewise, with respect to the reverse direction, in a *reverse stoichiometric matrix* $R \in \mathbb{Z}^{m\times n}$, let R_{ij} be the stoichiometry of molecule i participating as a substrate or catalyst in *reverse unidirectional reaction* j. The set of molecular species that jointly participate as either substrates or products in a single unidirectional reaction is referred to as a *reaction complex*.

One may define the topology of a *hypergraph* of reactions with a *net stoichiometric matrix* $S := R - F$. However, a catalyst, by definition, participates in a reaction with the same stoichiometry as a substrate or product ($F_{ij} = R_{ij}$), so the corresponding row of S is all zeros unless that catalyst is synthesised or consumed elsewhere in the same biochemical network, as is the case for many biochemical catalysts (Thiele et al., 2009). For example, consider the ith molecular species acting as a catalyst in some reactions. If it is synthesised in the jth reaction of a biochemical network, the stoichiometric coefficient in the forward stoichiometric matrix will be less than that of the forward stoichiometric matrix ($F_{ij} < R_{ij}$), so $S_{ij} := R_{ij} - F_{ij} > 0$. This also encompasses the case of an auto-catalytic reaction.

Before proceeding, some comments on our assumptions are in order. One may derive S from F and R, but the latter pair of matrices cannot, in general, be derived from S because S omits the stoichiometry of catalysis. The orientation of the hypergraph, i.e., the assignment of one direction to be forward (substrates $\rightharpoonup$ products), with the other reverse, is typically made so that net flux is forward (with positive sign) when a reaction is active in its biologically typical direction in

a biochemical network. This is an arbitrary convention rather than a constraint, and reversing the orientation of one reaction only exchanges one column of F for the corresponding one in R. Although every chemical reaction is in principle reversible, in a biochemical setting, due to physiological limits on the relative concentrations of reactants and substrates, some reactions are practically irreversible (Noor et al., 2013). Our conclusions also extend to systems of irreversible reactions because the reaction complexes for an irreversible reaction are the same as those for a reversible reaction.

In the following, the exponential or natural logarithm of a vector is meant component-wise, with $\exp(\log(0)) := 0$. Let $v_f \in \mathbb{R}^n_{>0}$ and $v_r \in \mathbb{R}^n_{>0}$ denote forward and reverse unidirectional reaction rate vectors. We assume that the rate of a unidirectional reaction is proportional to the product of the concentrations of each substrate or catalyst, each to the power of their respective stoichiometry in that unidirectional reaction (Wilhelmy, 1850), with linear proportionality given by strictly positive *rate coefficients* $k_f, k_r \in \mathbb{R}^n_{>0}$. Therefore we have

$$\begin{aligned} v_f(c) &:= \exp(\ln(k_f) + F^T \ln(c)), \\ v_r(c) &:= \exp(\ln(k_r) + R^T \ln(c)), \end{aligned} \tag{1}$$

where $c \in \mathbb{R}^m_{\geq 0}$ are molecular species concentrations. Strictly, it is not proper to take the logarithm of a unit that has physical dimensions, so c should be termed a vector of mole fractions rather than concentrations (Berry et al., 2000, Eq. 19.93), but safe in the knowledge that we have taken this liberty, we continue in terms of concentrations.

If the jth columns of F and R represent the stoichiometry of an *elementary reaction*, then the respective jth unidirectional reaction rate is given by an elementary kinetic rate law in (1). In biochemical modelling, often it is *composite reaction* stoichiometry that is represented, in which case the unidirectional reaction rates are given by pseudo-elementary kinetic rate laws. We shall revisit this point in discussion, but for now it suffices to mention that, in principle, all composite reactions can be decomposed into a set of elementary reactions following elementary reaction kinetics (Cook and Cleland, 2007), even allosteric reactions (Bray and Duke, 2004). With respect to the forward direction of an elementary reaction, the term *reaction complex* implies a corresponding physical association between substrate molecular species. For the sake of simplicity, we also use the term reaction complex for composite reactions, as if there were a corresponding simultaneous physical association of all substrates, which is generally not the case because composite reactions occur as a set of elementary reaction steps.

With respect to time, the deterministic rate of change of concentration is

$$\frac{dc}{dt} := (R - F)(v_f(c) - v_r(c)), \tag{2}$$

$$= ([R,\, F] - [F,\, R]) \begin{bmatrix} v_f(c) \\ v_r(c) \end{bmatrix}, \tag{3}$$

where $v_f(c) - v_r(c)$ gives a vector of net reaction rates, $[\cdot,\, \cdot]$ denotes the horizontal concatenation operator, and $:=$ denotes "is defined to be equal to". Time-invariant fluxes or concentrations satisfy (2) with $dc/dt := 0$. Define $k := \begin{bmatrix} k_f \\ k_r \end{bmatrix} \in \mathbb{R}^{2n}_{>0}$ to be given constants, then consider the *flux function*

$$v(c) := \exp(\ln(k) + [F,\, R]^T \ln(c)) = \begin{bmatrix} v_f(c) \\ v_r(c) \end{bmatrix} \tag{4}$$

with a concentration vector c the only argument. Apart from (a) our deliberate distinction between unidirectional and net stoichiometry, (b) our deliberate use of matrix-vector notation, and (c) our deliberate use of component-wise exponential and logarithm, the expression for unidirectional rate in (4) is a standard representation of deterministic elementary reaction kinetics.

2.2. Linear algebraic characterisation of flux-concentration duality

Herein, duality is defined as a one-to-one relationship between two variable vectors. We now establish a linear algebraic condition for duality between unidirectional flux and concentration vectors. This linear algebraic condition is a well known result in mathematics, but to our knowledge its application to establish duality between unidirectional flux and molecular species concentration is novel.

Theorem 1. *Assume we are given constants $k \in \mathbb{R}^{2n}_{>0}$ and $F, R \in \mathbb{Z}^{m\times n}_{\geq 0}$. Suppose a unidirectional reaction flux vector $v \in \mathbb{R}^{2n}_{>0}$ and a molecular species concentration vector $c \in \mathbb{R}^{m}_{>0}$ satisfy*

$$v = \exp(\ln(k) + [F,\, R]^T \ln(c)). \tag{5}$$

Then rank$([F,\, R]) = m$ is a necessary and sufficient condition for duality between fluxes and concentrations.

Proof. That v is uniquely defined given c is trivial. Taking the logarithm of both sides of (5), we have $\ln(v) - \ln(k) = [F,\, R]^T \ln(c)$. Then, if and only if $\text{rank}([F,\, R]) = m$ is $\ln(c)$, and therefore c, uniquely defined given v. □

Theorem 1 establishes that the flux function (4) is an injective function. It is not bijective because one can always find a v such that $\ln(v) - \ln(k)$ is not in the range of $[F,\, R]^T$. Note that the exponential function is bijective, but if one wished to consider other flux functions, it would be sufficient to replace the exponential function with another injective function and Theorem 1 would still hold.

We now proceed to interpret this stoichiometric condition for duality in biochemical terms. Consider the following triplet of isomerisation reactions involving three molecular species:

$$A \rightleftharpoons B, \quad B \rightleftharpoons C, \quad C \rightleftharpoons A.$$

The forward, reverse and net stoichiometric matrices are

$$F = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad R = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad (R - F) = \begin{bmatrix} -1 & 0 & 1 \\ 1 & -1 & 0 \\ 0 & 1 & -1 \end{bmatrix}, \tag{6}$$

where flux and concentration vectors are dual vectors because $\text{rank}([F,\, R]) = 3 = m$. Consider the following quartet of reactions involving four representatives of supposedly distinct molecular species:

$$A \rightleftharpoons B + C, \quad A \rightleftharpoons D, \quad B + C \rightleftharpoons D, \quad A + D \rightleftharpoons 2B + 2C.$$

The forward, reverse and net stoichiometric matrices are

$$F = \begin{bmatrix} 1 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad R = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 2 \\ 1 & 0 & 0 & 2 \\ 0 & 1 & 1 & 0 \end{bmatrix}, \quad (R - F) = \begin{bmatrix} -1 & -1 & 0 & -1 \\ 1 & 0 & -1 & 2 \\ 1 & 0 & -1 & 2 \\ 0 & 1 & 1 & -1 \end{bmatrix}, \tag{7}$$

where flux and concentration vectors are *not* dual vectors because rank($[F, R]$) = 3 < 4 = m. Observe that the second and third rows of F and R are positive multiples of one another. This corresponds to a pair of supposedly distinct molecules, B and C, that are always either produced or consumed together with fixed relative stoichiometry. This is an ambiguous model of reaction stoichiometry because either (i) B and C are actually the same molecular species and therefore the extra row is superfluous, or (ii) B and C are different molecular species but the model is missing some reaction that would demonstrate they are synthesised or consumed in distinct reactions.

2.3. Combinatorial characterisation of flux-concentration duality

The aforementioned linear algebraic condition for duality between unidirectional flux and concentration vectors is hard to interpret in terms of reaction complex stoichiometry. Therefore we sought a characterisation that would be easier to interpret in a (bio)chemically interpretable manner. Here we derive a combinatorial characterisation of the condition rank($[F, R]$) = m, which holds independently of the actual values of the stoichiometric coefficients. Our analysis draws from the theory of L-matrices and zero/sign patterns (Hershkowitz and Schneider, 1993; Brualdi and Shader, 2009). First we introduce some definitions and notation.

Definition 1. (Support of a set of vectors) Let C be a collection of d-dimensional row vectors. The support of C is defined to be the subset of $\mathcal{I} := \{1, \ldots, d\}$ such that, for each i in the given subset of $\mathcal{I}$, there exists at least one vector in C whose ith component is nonzero.

For example, if C is formed by the last two rows of the matrix

$$\begin{bmatrix} 1 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 \end{bmatrix},$$

the support of C is $\{2, 3, 5, 6\}$. If C is formed by the first and third columns of the matrix, its support is $\{1, 3\}$.

Definition 2. (Combinatorial independence) A collection C of row vectors (of equal dimension) is said to be combinatorially independent if C does not contain the zero vector and *every* two nonempty disjoint subsets of C have different supports.

In the above example, the rows of the matrix are combinatorially independent. However, the columns of this matrix are not combinatorially independent because the support of columns $\{1, 2\}$ is $\{1, 2, 3\}$, which is the same as the support of columns $\{3, 5\}$.

Definition 3. (Zero pattern) The zero pattern of a real matrix A is the $(0, 1)$-matrix obtained by replacing each nonzero entry of A by 1.

Theorem 2. *(Combinatorial independence and rank (Richman and Schneider, 1978, Lemma (5.2))) Let P be an $m \times d$ zero pattern. Every non-negative matrix with zero pattern P has rank m if and only if the rows of P are combinatorially independent.*

Conversely, it follows that if any two disjoint subsets of rows of P have the same support, then P is row rank-deficient. For example, the matrix

$$\begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \end{bmatrix}$$

is row rank-deficient because rows $\{1, 2, 3\}$ and rows $\{4, 5, 6\}$ have the same support $\{1, 2, 3, 4, 5, 6\}$. Theorem 2 permits us to state the following.

Theorem 3. *(Combinatorial independence and duality) Consider a family of biochemical networks that share the same zero pattern as* $[F, R]$*. Assume that each molecular species participates in at least one reaction in each network in the family. Then, for each network in the family, the following are equivalent:*

1. *The matrix* $[F, R]$ *has full row rank.*
2. *For every two disjoint sets of molecular species, there is at least one reaction complex that involves species from only one of the two sets.*
3. *Unidirectional flux and concentration are dual variables.*

Equivalently, for a given biochemical network with matrix $[F, R]$, if condition 2 in Theorem 3 is true, then one can exchange *any* positive stoichiometric coefficient of the network with *any* positive value and $[F, R]$ will still have full row rank. The above result provides a combinatorial characterisation of the condition for flux-concentration duality, which holds independent of the values of the stoichiometric coefficients. This is analogous to results involving L-matrices for problems such as the structural controllability of systems (Brualdi and Shader, 2009).

2.3.1. Testing for combinatorial independence

According to Theorem 2, to test if an $m \times d$ zero pattern has rank m, we can equivalently test whether its m rows are combinatorially independent. Can this test be performed efficiently? In general (unless P=NP) the answer is no (Klee et al., 1984), as the problem of testing if a *sign pattern* (elements $\{0, 1, -1\}$) has full row rank is NP-complete. Their proof relies on a reduction from the 3-SAT problem, which is known to be NP-complete (Garey and Johnson, 1979). In their proof they construct a *non-negative* sign pattern (which is a zero pattern), and therefore their result applies to our case too. Hence we have the following.

Theorem 4. *(Testing combinatorial independence* ***(Klee et al., 1984)****) Let P be a zero pattern. Testing if the rows of P are combinatorially independent is NP-complete.*

However, as we prove next, when the zero pattern is constrained to have at most two non-negative entries per column, the testing for combinatorial independence can be done in polynomial time. To our knowledge, this result is new.

Theorem 5. *(Testing combinatorial independence in constrained zero patterns) Let P be a zero pattern with at most two 1s per column. Testing if the rows of P are combinatorially independent can be done in polynomial time.*

Proof. Without loss of generality we can assume that each column of P has exactly two nonzero entries. We view the matrix P as the incidence matrix of an undirected graph, where each row of P is a vertex and each column is an edge. Combinatorial dependence of the rows of P would imply the existence of two disjoint sets of rows with the same support, which would imply the existence of a connected component of the graph that is bipartite (2-colorable). Finding all connected components of a graph and bipartiteness testing are classical graph problems that can be solved in polynomial time (Cormen et al., 2009). □

Since most reconstructed biochemical networks are in terms of composite reactions, the corresponding $[F, R]$ may have more than two nonzero entries per column and the nonzero stoichiometric coefficients may differ from 1. However, every composite reaction is a composition of a set of elementary reactions (Cook and Cleland, 2007), each with at most three reactants per reaction, so the resulting *bilinear* $[F, R]$ will have at most two nonzero entries per column. It is possible to algorithmically convert any composite reaction into a set of elementary reactions, with at most two nonzero entries per column, by creating faux molecular species representing a reaction intermediate, e.g., the composite reaction $A + B \rightleftharpoons C + D$ may be decomposed into $A + B \rightleftharpoons E$ and $E \rightleftharpoons C + D$. Reaction intermediates are typically not identical for two enzyme-catalysed composite reactions, suggesting that flux-concentration duality is a pervasive property of biochemical networks in general.

2.4. Flux-concentration duality in existing genome-scale biochemical networks

Section 2.3 provided a biochemically interpretable condition, in terms of molecular species involvement in reaction complex stoichiometry, that implies flux-concentration duality for an arbitrary network. We now show that flux-concentration duality is a pervasive property of quality-controlled models derived from genome-scale biochemical network reconstructions. Testing for combinatorial independence is computationally complex, so instead we rely on linear algebra to test the rank of $[F, R]$. As detailed below, we converted 29 genome-scale metabolic network reconstructions into computational models, then compared the number of molecular species with the rank of $[F, R]$ before and after conversion. These metabolic reconstructions were all manually curated and represent a wide range of different species (see Supplementary Table 1).

It is important to distinguish a network reconstruction from a computational model of a biochemical network. The former may contain incomplete or inconsistent knowledge of biochemistry, while the latter must satisfy certain modelling assumptions, represented by mathematical conditions, in order to ensure that the model is a faithful representation of the underlying biochemistry. This modelling principle is already well established in the digital circuit modelling community, and some of the associated *model checking* algorithms have been applied to biochemical networks (Carrillo et al., 2012), especially by the community that use Petri-nets to model biochemical networks, e.g., (Soliman, 2012). The application of modelling assumptions is a key step in the conversion of a reconstruction into a computational model. We now introduce these assumptions, their mathematical representation, and their relationship to the rank of $[F, R]$. For the sake of simplicity, the toy examples given to illustrate key concepts only involve reactions with two or less reactants, but the theory presented also applies to systems of composite reactions involving three or more reactants.

2.4.1. Stoichiometric consistency

All biochemical reactions conserve mass; therefore it is essential in a model that each reaction, which is supposed to represent a biochemical reaction, does actually conserve mass.

Although it is not essential to do so (Fleming and Thiele, 2012), reactions that do not conserve mass are often added to a network reconstruction (Thiele and Palsson, 2010) in order to represent the flow of mass into and out of a system, e.g., during flux balance analysis (Palsson, 2006). Every reaction that does not conserve mass, but is added to a model in order represent the exchange of mass across the boundary of a biochemical system, is henceforth referred to as an *exchange reaction*, e.g., $D \rightleftharpoons \emptyset$, where $\emptyset$ represents null. When checking for reactions that do not conserve mass, we must first omit exchange reactions.

Besides exchange reactions, a reconstruction may contain reactions with incompletely specified stoichiometry or molecules with incompletely specified chemical formulae, because of (for instance) limitations in the available literature evidence. While stoichiometrically inconsistent biochemical reactions may appear in a reconstruction, they should be omitted from a computational model derived from that reconstruction, especially if the model is to be used to predict flow of mass, else erroneous predictions could result. One approach is to require that chemical formulae be collected for each molecule during the reconstruction process (Thorleifsson and Thiele, 2011), then omit non-exchange reactions that are elementally imbalanced (Schellenberger et al., 2011). A complementary approach is to detect reactions that are specified in a *stoichiometrically inconsistent* manner (Gevorgyan et al., 2008). For instance, the reactions $A + B \rightleftharpoons C$ and $C \rightleftharpoons A$ are stoichiometrically inconsistent because it is impossible to assign a positive molecular mass to all species whilst ensuring that each reaction conserves mass.

A set of stoichiometrically consistent reactions is mathematically defined by the existence of at least one $\ell \in \mathbb{R}^m_{>0}$ such that $R^T\ell = F^T\ell$, equivalently $S^T\ell = (R-F)^T\ell = 0$, where ℓ is a vector of the molecular mass of m molecular species. Consider the aforementioned stoichiometrically inconsistent example, where the corresponding stoichiometric matrices are

$$S := R - F = \begin{bmatrix} 0 & 1 \\ 0 & 0 \\ 1 & 0 \end{bmatrix} - \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} -1 & 1 \\ -1 & 0 \\ 1 & -1 \end{bmatrix},$$

with rows from top to bottom corresponding to molecular species A, B, C. Let $a, b, c \in \mathbb{R}$ denote the molecular mass of A, B, C. We require a, b, c such that

$$R^T\ell = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} c \\ a \end{bmatrix} = \begin{bmatrix} a+b \\ c \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = F^T\ell.$$

However, the only solution requires $a = c$ and $b = 0$, i.e., a zero mass for the molecule B, which is inconsistent with chemistry; therefore the reactions $A + B \rightleftharpoons C$ and $C \rightleftharpoons A$ are stoichiometrically inconsistent. In general, given F and R, one may check for stoichiometric consistency (Gevorgyan et al., 2008) by solving the optimisation problem

$$\begin{aligned} \max_{\ell} \ & \|\ell\|_0 \\ \text{s.t. } & S^T\ell = 0, \\ & 0 \leq \ell. \end{aligned}$$

Here, $\|\ell\|_0$ denotes the zero-norm or equivalently the cardinality (number of non-zero entries) of ℓ. However, maximising the cardinality of a non-negative vector in the left nullspace of S is a problem that is challenging to solve exactly. This problem has been represented as a mixed-integer linear optimisation problem (Gevorgyan et al., 2008), but since algorithms for such prob-

lems have unpredictable computational complexity, we implemented a novel and more efficient approach.

The cardinality of a non-negative vector is a quasiconcave (or unimodal) function (Boyd and Vandenberghe, 2004). The problem of maximising this particular quasiconcave function, subject to a convex constraint, may be approximated by a linear optimisation problem (Vlassis et al., 2014), in our case the problem

$$\begin{aligned} \max_{z,\ell} \ & \mathbb{1}^T z \\ \text{s.t. } & S^T \ell = 0, \\ & z \leq \ell, \\ & 0 \leq z \leq \mathbb{1}\alpha, \\ & 0 \leq \ell \leq \mathbb{1}\beta, \end{aligned} \tag{8}$$

where $z, \ell \in \mathbb{R}^m$ and $\mathbb{1}$ denotes an all ones vector. In this approximation, we maximise the sum over all dummy variables z_i, $i = 1, \ldots, m$, but it is ℓ_i that represents the stoichiometrically consistent molecular mass of the ith molecule. The scalars $\alpha, \beta \in \mathbb{R} > 0$ are proportional to the smallest molecular mass considered non-zero and the largest molecular mass allowed. An upper bound on the largest molecular mass avoids the possibility of a poorly scaled optimal ℓ. We used $\alpha = 10^{-4}$ and $\beta = 10^4$ as all models tested were of metabolism, so eight orders of magnitude between the least and most massive metabolite is sufficient. As this approximation is based on linear optimisation, it can be implemented numerically in a scalable manner. We applied (8) to each reconstruction in Supplementary Table 1 in order to identify stoichiometrically inconsistent rows. That is, if $\ell^\star$ denotes the optimal ℓ obtained from (8) then the ith row is stoichiometrically inconsistent if $\ell_i^\star < \alpha$. Stoichiometrically inconsistent rows and the corresponding columns were omitted from further analyses. Where molecular formulae were available, we confirmed that all retained biochemical reactions were elementally balanced, as expected. To reiterate, in our numerical check of rank $[F, R]$, discussed below, all rows correspond to metabolite species involved in stoichiometrically consistent reactions, with the exception of exchange reactions.

2.4.2. Net flux consistency

If one assumes that all molecules are at steady state, the corresponding computational model should be *net flux consistent*, meaning that each net reaction of the network has a nonzero flux in at least one feasible steady state net flux vector. Due to incomplete biochemical knowledge, a reconstruction may contain *net flux inconsistent* reactions that do not admit a nonzero steady state net flux. For example, consider the set of reactions

$$\emptyset \rightleftharpoons D \rightleftharpoons G \rightleftharpoons \emptyset, \qquad D \rightleftharpoons H. \tag{9}$$

In this set, the reaction $D \rightleftharpoons H$ is net flux inconsistent, as any nonzero net flux is inconsistent with the assumption that the concentration of C should be time invariant. Inclusion of net flux inconsistent reactions, like $D \rightleftharpoons H$, in a dynamic model would be perfectly reasonable, but we omit such reactions because the focus of this paper is on modelling of steady states.

Let $B \in \mathbb{R}^{m \times p}$ denote the stoichiometric matrix for a set of p exchange reactions. We say a matrix S is net flux consistent if there exist matrices $V \in \mathbb{R}^{n \times k}$ and $W \in \mathbb{R}^{p \times k}$ such that

$$SV = -BW,$$

where each row of V and each row of W contains at least one nonzero entry. Consider the aforementioned net flux inconsistent example, where the corresponding stoichiometric matrices are

$$S = \begin{bmatrix} -1 & -1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \qquad B = \begin{bmatrix} 1 & 0 \\ 0 & -1 \\ 0 & 0 \end{bmatrix}.$$

Let $p, q, r, s \in \mathbb{R}$ denote the net rate of the reactions, from left to right in (9). We require p, q, r, s such that

$$SV = \begin{bmatrix} -1 & -1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} p \\ q \end{bmatrix} = \begin{bmatrix} -p-q \\ q \\ q \end{bmatrix} = \begin{bmatrix} -r \\ s \\ 0 \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} r \\ s \end{bmatrix} = -BW.$$

However, the only solution requires $q = 0$, i.e., a zero net flux through the reaction $D \rightleftharpoons H$, corresponding to a zero row of V; therefore this reaction is net flux inconsistent. Our definition of net flux consistency is weaker than the assumption that all reactions admit a nonzero net flux simultaneously, which would be equivalent to requiring a single net flux vector with all nonzero entries, i.e., $k = 1$. It is also weaker than the assumption of net flux consistency subject to bounds on the direction of reactions (Vlassis et al., 2014), which we do not impose here. Enforcing net flux consistency requires omission of any net reaction that cannot carry a non-zero net flux at a steady state.

Within fastCORE, a scalable algorithm for reconstruction of compact and context-specific biochemical network models (Vlassis et al., 2014), a key step employs linear optimisation as described above (8) to identify the largest set of net flux consistent reactions in a given model. We created a computational model from the stoichiometrically consistent subset of each reconstruction in Supplementary Table 1. We allowed all reactions to be reversible (lower and upper bounds -1000 and 1000), included exchange reactions in each reconstruction, and then identified and omitted all net flux inconsistent reactions ($v_j < \epsilon = 10^{-4}$). We also omitted the corresponding rows, where a molecular species is only involved in flux inconsistent reactions. Therefore, in our check of rank($[F, R]$), all rows correspond to metabolite species involved in net flux consistent reactions. As Supplementary Table 1 illustrates, this is typically a subset of the stoichiometrically consistent rows.

2.4.3. Unique and non-trivial molecular species

In a reconstruction, one may find a pair of rows in S that are identical up to scalar multiplication. As these extra rows typically represent inadvertent duplication of an identical molecular species, any such duplicate rows were omitted. Likewise, we omitted any row with all zeros, e.g., corresponding to a metabolite that was only involved in stoichiometrically inconsistent or net flux inconsistent reactions. Hereafter, any biochemical network without zero rows or rows identical up to scalar multiplication we refer to as being *non-trivial*.

2.4.4. Pervasive flux-concentration duality

We investigated the stoichiometric properties of a representative subset of published metabolic network reconstructions. Specifically, numerical experiments were performed on 29 published reconstructions where a Systems Biology Markup Language (Keating et al., 2006) compliant Extensible Markup Language (.xml) file was available and at least 90% of the molecular species corresponded to stoichiometrically consistent rows. Numerical linear algebra was used to compute

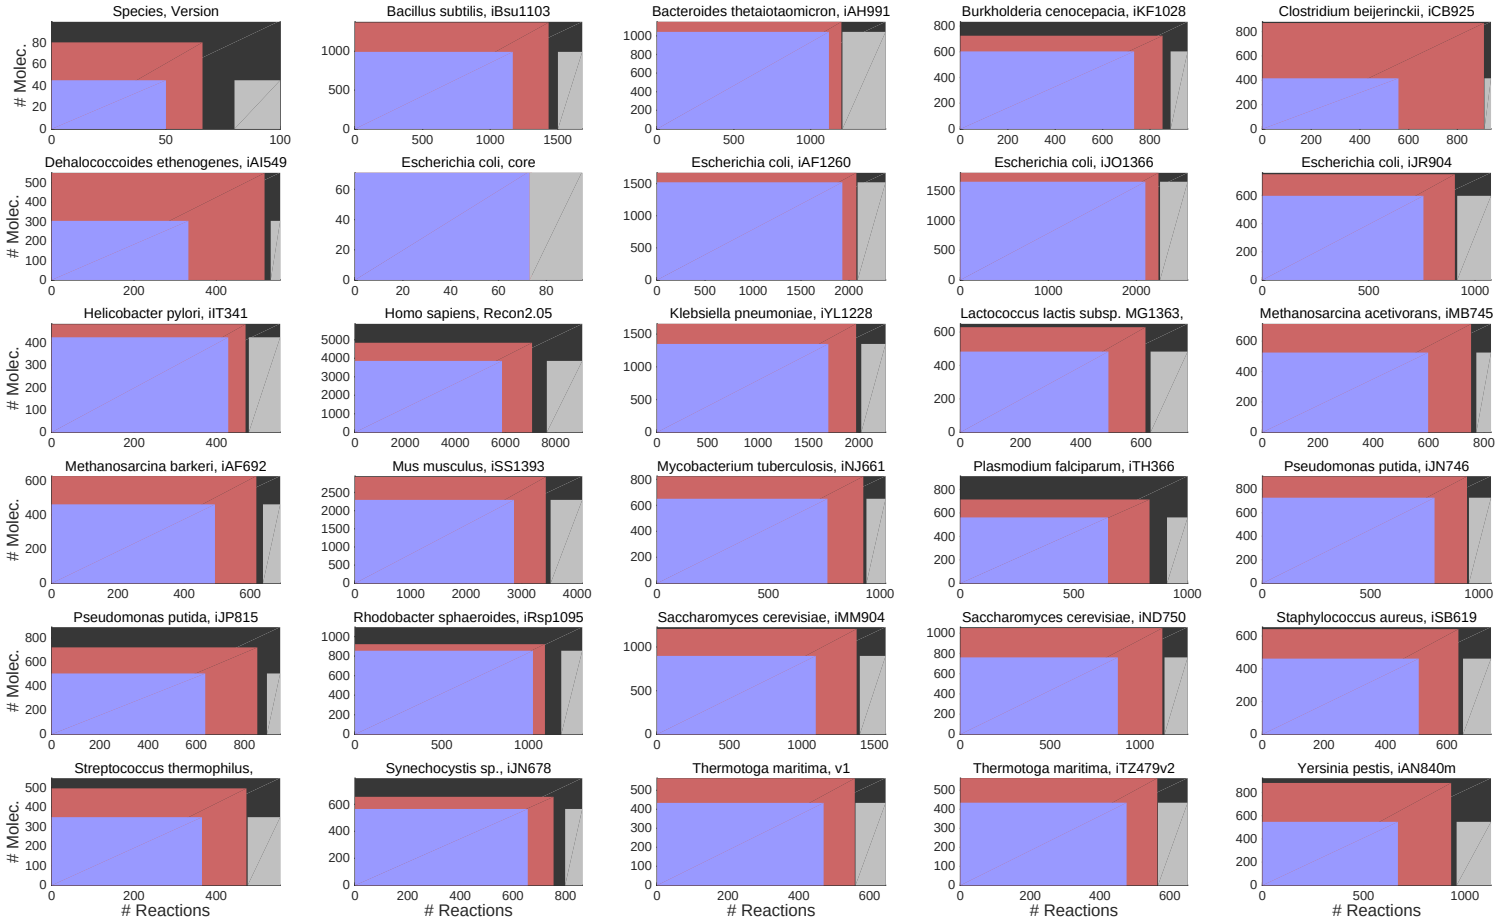

Figure 1: Usually, only a subset of a reconstruction will satisfy the mathematical conditions imposed when a corresponding computational model is generated. The original size of $[S,\ S_e]$ (outer black rectangle) varies across the 29 reconstructions tested. Due to exchange reactions, only a subset of the columns of a reconstruction correspond to stoichiometrically consistent rows (red rectangles). If a molecular species is exclusively involved in exchange reactions, the number of stoichiometrically consistent rows is less than the number of rows of reconstruction. Due to reactions that do not admit a nonzero steady state net flux, only a subset of mass balanced reactions and a subset of exchange reactions are also flux consistent (blue and grey rectangles, respectively). When F and R are derived from a subset of a genome-scale biochemical network reconstruction, assuming no zero rows of $[F,\ R]$ and no rows that are identical up to scalar multiplication, stoichiometric and net flux consistency is often but not always sufficient to ensure that $[F,\ R]$ has full row rank.

matrix rank (cf. Supplementary File 1, Section 5.1.1). The results are summarised in Figure 1 and provided in detail in Supplementary File 2. All numerical experiments may be reproduced with the MATLAB code distributed with the COBRA Toolbox at `https://github.com/opencobra/cobratoolbox` (cf. Supplementary File 1, Section 5.3).

The number of (possibly indistinct) molecular species is, by definition, equivalent to the number of rows of $S := R - F$ derived directly from the reconstruction, without additional assumptions. By forming $[F, R]$ directly from a reconstruction, we found that rank($[F, R]$) is usually (21/29) less than the number of rows of S, with some (8/29) exceptions, e.g., the genome-scale reconstruction of the metabolic network of *Rhodobacter sphaeroides*, iRsp1095 (Imam et al., 2011).

Most genome-scale reconstructions (26/29) were accompanied by chemical formulae for the majority of reactions. If the number of stoichiometrically consistent rows is less than the number of molecules exclusively involved in reactions that are supposed to be elementally balanced, as determined by a check for elemental balance, then at least one chemical formula for a molecular species must be incorrectly specified. In only 3 of the 26 reconstructions that supplied chemical formulae, this issue was apparent (cf. Supplementary file 1). Each reconstruction was converted into a computational model where $F, R \in \mathbb{R}_{\geq 0}^{m \times n}$ satisfy the following conditions:

1. All rows of $S := R - F$ correspond to molecular species in stoichiometrically consistent reactions, with the exception of exchange reactions.
2. No two rows in $[F, R]$ are identical up to scalar multiplication.
3. All rows of S correspond to molecular species in net flux consistent reactions, assuming all reactions are reversible, including exchange reactions.
4. No row of $[F, R]$ is all zeros.

Of the 29 reconstructions subjected to the aforementioned conditions, 26 generated a model where $[F, R]$ had full row rank. When $[F, R]$ was row rank-deficient, the rank was never more than three less than the number of rows of $[F, R]$. In each case, the rank-deficiency was a result of omitted biochemical reactions that would otherwise have resulted in an $[F, R]$ with full row rank. A typical example of a genome-scale reconstruction with row rank-deficient $[F, R]$ is highlighted in Section 5.2. In general, should a row rank-deficient $[F, R]$ arise, there are two options: (i) further manual reconstruction effort to correctly specify reaction network stoichiometry, or (ii) omission of the dependent molecular species from any derived kinetic model.

Although conditions 2 and 4 are trivial and clearly necessary, neither of conditions 1 or 3 (stoichiometric consistency or net flux consistency) is necessary for $[F, R]$ to have full row rank. For almost one third (8/29) of the reconstructions, one could form $[F, R]$ without any further assumptions and yet $[F, R]$ had full row rank. For instance, the genome-scale *Methanosarcina acetivorans C2A* metabolic model (iMB745 (Benedict et al., 2012)) has 715 molecular species and without stoichiometric or net flux consistency being imposed, rank($[F, R]$) = 715, even though this is 2 greater than the number of stoichiometrically consistent rows of S.

When a stoichiometrically inconsistent row of S is omitted from a metabolic model, the corresponding row of the biomass reaction is also omitted. This reduction in the number of constraints could lead to an increase in the maximum biomass synthesis rate. In contrast, removal of net flux inconsistent reactions might reduce the maximum biomass synthesis rate or render biomass synthesis infeasible. Flux balance analysis of each of the 29 genome-scale reconstructions before and after application of the aforementioned four conditions revealed that growth feasibility was not extinguished and tended to increase (data not shown). Further iterations of reconstruction and model validation would be required for each model derived in the manner

described above prior to use in applications. In particular, one should check that each omitted reaction is balanced for each atomic element and conduct further literature research to resolve flux inconsistent reactions that contributed toward optimal biomass synthesis in models derived from reconstructions without the aforementioned quality control steps.

3. Discussion

Any net stoichiometric matrix $S \in \mathbb{R}^{m\times n}$ may be derived by taking the difference between a pair of forward and reverse stoichiometric matrices $F, R \in \mathbb{R}^{m\times n}_{\geq 0}$, that is $S := R - F$. The horizontal concatenation $[F,\ R] \in \mathbb{R}^{m\times 2n}$ is a key mathematical object that appears in the deterministic, elementary, unidirectional reaction kinetic rate equation $v = \exp(\ln(k) + [F,\ R]^T \ln(c))$, relating concentrations $c \in \mathbb{R}^m$ and rate coefficients $k \in \mathbb{R}^{2n}$ to fluxes $v \in \mathbb{R}^{2n}$. We address the question: When does there exist a one-to-one relationship between concentrations and fluxes?

We have proven that, given rate coefficients, there is a one-to-one relationship between concentrations and fluxes if and only if $[F,\ R]$ has full row rank. Furthermore, this dual relationship exists if and only if there are no two disjoint sets of molecular species where every corresponding unidirectional reaction involves at least one molecular species from each of the disjoint sets. Flux-concentration duality implies that one could discuss biochemistry either entirely in terms of fluxes or entirely in terms of concentrations, as both would be different perspectives on the same biochemical system. This has clear implications when interpreting biochemical network function from the perspective of either concentrations or fluxes.

Within a wide range of non-trivial biochemical network reconstructions, including metabolism and signalling networks, we observe from numerical experiments that together, stoichiometric and net flux consistency of S is often sufficient to ensure that $[F,\ R]$ has full row rank. After application of these conditions we occasionally observe that $[F,\ R]$ is row rank-deficient and this is due to omission of reactions from the corresponding reconstruction. Finding a numerical example where $[F,\ R]$ is row rank-deficient does not reduce the biochemical significance of our observations if the underlying network is not biochemically realistic. In each particular case, it was clear that row rank-deficiency $[F,\ R]$ was due to the omission of known biochemical reactions that would have given $[F,\ R]$ full row rank. It is easy to test if $[F,\ R]$ has full row rank for a particular network, but it is a rather abstract linear algebraic condition, so it is not easy to see if it applies to biochemical networks in general. Therefore, we sought a complementary characterisation of full-row-rank $[F,\ R]$ that was applicable in general and more easily interpretable from a biochemical network perspective.

We have established biochemically interpretable combinatorial conditions that are necessary and sufficient for $[F,\ R]$ to have full row rank dependent only on the sparsity pattern of F and R; that is, independent of the actual values of their nonzero entries. However, in practice these combinatorial conditions may be too strong, because for any given biochemical network, the values of the nonzero entries are fixed and the corresponding $[F,\ R]$ may have full row rank, even if combinatorial independence of its rows does not hold. Combinatorial independence of the rows of a given $[F,\ R]$ implies full row rank, but in general, the reverse implication does not hold. In Section 2.4, we applied numerical linear algebra to check the rank of $[F,\ R]$ derived from 29 reconstructions, each subject to certain conditions. However, as the aforementioned $[F,\ R]$ all correspond to networks of composite biochemical reactions, there exist columns of $[F,\ R]$ with more than two nonzero entries. We do not test for combinatorial independence of the rows of these $[F,\ R]$, as this problem is NP-hard (Garey and Johnson, 1979).

There are many interesting open problems, the solution of which would be interesting extensions to this work. We know that all composite reactions are defined from the composition of a set of elementary reactions, and the latter give rise to an $[F, R]$ with at most two nonzero entries in each column. Given an $[F, R]$ derived from a network of composite reactions, if one were to express the network as a set of elementary reactions that properly reflects the underlying biochemistry (Cook and Cleland, 2007), does the corresponding $[F, R]$ also have full row rank? One could ask the same question starting from an elementary reaction network with an $[F, R]$ that has full row rank. Indeed, by Theorem 4, testing the combinatorial independence of the latter is solvable in polynomial time. It is exciting that so many of the non-trivial, stoichiometrically consistent and net flux consistent biochemical networks that we tested do give rise to an $[F, R]$ of full row rank, despite the fact that mathematically we know that these conditions are not sufficient for $[F, R]$ to have full row rank. What are the undiscovered, necessary, mathematical, yet biologically interpretable conditions that ensure $[F, R]$ has full row rank, even if its rows are not combinatorially independent?

Putting this work into a broader context, one must always make a clear distinction between a reconstruction and a model. In practice, the latter is a numerical implementation that must satisfy certain mathematical conditions that are usually not satisfied by every metabolite species and every reaction in a given reconstruction. Indeed, depending on one's combination of mathematical assumptions, one could derive many different models from the same reconstruction. Testing for compliance with mathematical conditions is a vital element of quality control when converting a reconstruction into a correctly specified computational model. Of note in this respect is the relatively low computational complexity of the linear optimisation algorithms we use to solve the problem of checking for stoichiometric and net flux consistency.

Reconstruction mis-specification is often not due to some error, especially for reconstructions that are ambitious in scope. Such reconstructions will inevitably contain knowledge gaps, where the exact stoichiometry, chemical formula, etc, is unknown for certain reactions. That is, reconstruction mis-specification is often a reflection of incomplete biochemical knowledge. As any computational model will only represent the subset of the metabolite species and reactions that satisfy certain mathematical conditions, e.g., stoichiometric consistency, one must take care to omit that part of a reconstruction not satisfying certain conditions before generating model predictions and absolutely before making any biological conclusions. Otherwise grossly erroneous conclusions may be obtained.

In applied mathematics, the development of an algorithm to find a solution to a system of equations begins with certain assumptions on the properties of the function(s) involved. In systems biochemistry, deterministic modelling of molecular species concentrations gives rise to systems of nonlinear equations, e.g., (2), the general mathematical properties of which are still being discovered. Given rate coefficients, there is a paucity of scalable algorithms, with guaranteed convergence properties, to solve large nonlinear biochemical reaction equation systems for non-equilibrium, stationary concentrations. Likewise for the problem of fitting optimal rate coefficients given concentrations and a known reaction equation system. Observe that (2) contains the matrix $[F, R]$ twice and the matrix $[R, F]$ once.

That $\text{rank}([F, R]) = \text{rank}([R, F]) = m$ is a pervasive property of biochemical networks from a diverse set of organisms motivates the development of algorithms to exploit this property and its consequences, e.g., (Artacho et al., 2015). This algorithmic development proceeds with two complementary approaches: theory and numerical experiments. Of particular importance in this regard is that the set of models generated herein (with $\text{rank}([R, F]) = m$) satisfy a common set of mathematical conditions, thereby reducing the possibility for spurious numerical results, when

numerically testing hypothesised but unproven theorems concerning the properties of biochemical networks in general. For instance it is known that a full row rank $[R, F]$ is a necessary but insufficient condition to preclude the existence of multiple positive steady states for certain chemical reaction networks (Müller et al., 2015). Testing the rank of $[R, F]$ can be done efficiently, but it is still an open problem to design a tractable algorithm to test for the necessary and sufficient conditions to preclude the existence of multiple positive steady states for genome-scale biochemical networks (Müller et al., 2015). Numerical tests of a mathematical conjecture, using biochemically realistic stoichiometric matrices, can be an efficient way to find a counter-example or to provide support for the plausibility of a conjecture. These tests help one decide where to invest the mental effort required to attempt a proof of a conjecture. It is important therefore that such numerical tests be conducted with (a) a wide selection of stoichiometric matrices, in case a conjecture holds only for certain network topologies, and (b) a set of stoichiometric matrices that each satisfy a specified set of biochemically motivated mathematical conditions, in case a conjecture holds only for stoichiometric matrices corresponding to realistic biochemical networks.

4. Conclusions

Mathematical and computational modelling of biochemical networks is often done in terms of either the concentrations of molecular species or the fluxes of biochemical reactions. Mathematical modelling from either perspective is equivalent when concentrations and unidirectional fluxes are dual variables. Assuming elementary kinetic rate laws for each reaction, we show that this duality holds if and only if the matrix $[F, R] \in \mathbb{R}^{m\times 2n}_{\geq 0}$ has full row rank, where $[F, R]$ is formed by concatenation of the stoichiometric matrices $F, R \in \mathbb{R}^{m\times n}_{\geq 0}$ for the m reactants consumed in n forward and reverse reaction directions, respectively. Numerical experiments with computational models derived from many genome-scale biochemical networks indicate that flux-concentration duality is a pervasive property of biochemical networks. For an arbitrary biochemical network, we provide a combinatorial characterisation that is sufficient to ensure flux-concentration duality. That is, for every two disjoint sets of molecular species, if there is at least one reaction complex that involves species from only one of the two sets, then duality holds. Our stoichiometric characterisation of the conditions for duality between concentrations and unidirectional fluxes has fundamental implications for mathematical and computational modelling of biochemical networks. When flux-concentration duality holds, interpretation of biochemical network function from the perspective of unidirectional fluxes is equivalent to interpretation from the perspective of molecular species concentrations.

Acknowledgements

We would like to thank Michael Tsatsomeros and Francisco J. Aragon Artacho for valuable comments. This work was funded by the Interagency Modeling and Analysis Group, Multiscale Modeling Consortium U01 awards from the National Institute of General Medical Sciences [award GM102098] and U.S. Department of Energy, Office of Science, Biological and Environmental Research Program [award de-sc0010429]. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies.

References

References

Artacho, F. J. A., Fleming, R. M. T., Vuong, P. T., Jul. 2015. Accelerating the DC algorithm for smooth functions. arXiv:1507.07375 [math, q-bio]ArXiv: 1507.07375.

Ballerstein, K., Kamp, A. v., Klamt, S., Haus, U.-U., Feb. 2012. Minimal cut sets in a metabolic network are elementary modes in a dual network. Bioinformatics 28 (3), 381–387.

Benedict, M. N., Gonnerman, M. C., Metcalf, W. W., Price, N. D., Feb 2012. Genome-scale metabolic reconstruction and hypothesis testing in the methanogenic archaeon Methanosarcina acetivorans C2A. J Bacteriol 194 (4), 855–865.

Berry, S. R., Rice, S. A., Ross, J., 2000. Physical Chemistry, 2nd Edition. Oxford University Press, Oxford.

Boyd, S. P., Vandenberghe, L., 2004. Convex optimization. Cambridge University Press, Cambridge, UK; New York.

Bray, D., Duke, T., 2004. Conformational spread: the propagation of allosteric states in large multiprotein complexes. Annual review of biophysics and biomolecular structure 33, 53–73.

Brualdi, R. A., Shader, B. L., 2009. Matrices of Sign-solvable Linear Systems. Vol. 116. Cambridge University Press.

Carrillo, M., Góngora, P. A., Rosenblueth, D. A., Jul. 2012. An overview of existing modeling tools making use of model checking in the analysis of biochemical networks. Front Plant Sci 3:155.

Cook, P. F., Cleland, W. W., 2007. Enzyme Kinetics and Mechanism. Taylor & Francis Group, London.

Cormen, T. H., Leiserson, C. E., Rivest, R. L., Stein, C., 2009. Introduction to Algorithms, 3rd Edition. MIT press.

Fleming, R. M. T., Thiele, I., 2012. Mass conserved elementary kinetics is sufficient for the existence of a non-equilibrium steady state concentration. Journal of Theoretical Biology 314, 173–181.

Garey, M. R., Johnson, D. S., 1979. Computers and Intractability: a Guide to NP-completeness. WH Freeman, New York.

Gevorgyan, A., Poolman, M. G., Fell, D. A., 2008. Detection of stoichiometric inconsistencies in biomolecular models. Bioinformatics 24 (19), 2245–2251.

Gill, P. E., Murray, W., Saunders, M. A., Jan. 2005. SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization. SIAM Rev. 47 (1), 99–131.

Giovannoni, S. J., Tripp, H. J., Givan, S., Podar, M., Vergin, K. L., Baptista, D., Bibbs, L., Eads, J., Richardson, T. H., Noordewier, M., Rappé, M. S., Short, J. M., Carrington, J. C., Mathur, E. J., Aug. 2005. Genome streamlining in a cosmopolitan oceanic bacterium. Science 309 (5738), 1242–1245.

Haus, U.-U., Klamt, S., Stephen, T., Mar. 2008. Computing knock-out strategies in metabolic networks. Journal of Computational Biology 15 (3), 259–268.

Heinrich, R., Rapopoort, S. M., Rapoport, T. A., 1978. Metabolic regulation and mathematical models. Progress in Biophysics and Molecular Biology, 32 (1), 1–82.

Hershkowitz, D., Schneider, H., 1993. Ranks of zero patterns and sign patterns. Linear and Multilinear Algebra 34 (1), 3–19.

Imam, S., Yilmaz, S., Sohmen, U., Gorzalski, A., Reed, J., Noguera, D., Donohue, T., 2011. iRSP1095: A genome-scale reconstruction of the Rhodobacter sphaeroides metabolic network. BMC Systems Biology 5 (1), 116.

Jamshidi, N., Palsson, B. Ø., Feb. 2009. Flux-concentration duality in dynamic nonequilibrium biological networks. Biophysical Journal 97 (5), 11–13.

Keating, S. M., Bornstein, B. J., Finney, A., Hucka, M., May 2006. SBMLtoolbox: an SBML toolbox for matlab users. Bioinformatics 22 (10), 1275–1277.

Klamt, S., Gilles, E. D., Jan. 2004. Minimal cut sets in biochemical reaction networks. Bioinformatics 20 (2), 226–234.

Klee, V., Ladner, R., Manber, R., 1984. Sign solvability revisited. Linear Algebra and its Applications 59, 131–157.

Lewis, G., 1925. A new principle of equilibrium. Proc Natl Acad Sci USA 11 (3), 179–183.

Müller, S., Feliu, E., Regensburger, G., Conradi, C., and Shiu, A., Dickenstein, A., Oct. 2014. Sign conditions for injectivity of generalized polynomial maps with applications to chemical reaction networks and real algebraic geometry. Found Comput Math 1–29 doi:10.1007/s10208-014-9239-3

Mahadevan, R., von Kamp, A., Klamt, S., Apr. 2015. Genome-scale strain designs based on regulatory minimal cut sets. Bioinformatics, 31 (17), 2844–2851.

Noor, E., Haraldsdóttir, H. S., Milo, R., Fleming, R. M. T., Jul. 2013. Consistent estimation of Gibbs energy using component contributions. PLoS Comput Biol 9 (7), e1003098.

Palsson, B. Ø., 2006. Systems Biology: Properties of Reconstructed Networks. Cambridge University Press, Cambridge.

Richman, D. J., Schneider, H., 1978. On the singular graph and the Weyr characteristic of an M-matrix. Aequationes Mathematicae 17 (1), 208–234.

Schellenberger, J., Que, R., Fleming, R. M. T., Thiele, I., Orth, J. D., Feist, A. M., Zielinski, D. C., Bordbar, A., Lewis, N. E., Rahmanian, S., Kang, J., Hyduk, D., Palsson, B. Ø., 2011. Quantitative prediction of cellular metabolism with constraint-based models: the COBRA Toolbox v2.0. Nature Protocols 6 (9), 1290–1307.

Schuster, S., Fell, D. A., Dandekar, T., Mar. 2000. A general definition of metabolic pathways useful for systematic organization and analysis of complex metabolic networks. Nat Biotech 18 (3), 326–332.

Soliman, S., 2012. Invariants and other structural properties of biochemical models as a constraint satisfaction problem. Algorithms for Molecular Biology 7 (1), 15.

Sun, Y., Fleming, R. M. T., Thiele, I., Saunders, M. A., 2013. Robust flux balance analysis of multiscale biochemical reaction networks. BMC Bioinformatics 14 (1), 240.

Thiele, I., Fleming, R. M. T., Que, R., Bordbar, A., Diep, D., Palsson, B. Ø., 2012. Multiscale modeling of metabolism and macromolecular synthesis in *E. coli* and its application to the evolution of codon usage. PLoS One 7 (9), e45635.

Thiele, I., Jamshidi, N., Fleming, R. M. T., Palsson, B. Ø., 2009. Genome-scale reconstruction of *E. coli*'s transcriptional and translational machinery: a knowledge-base, its mathematical formulation, and its functional characterization. PLoS Comp Biol 5 (3), e1000312.

Thiele, I., Palsson, B. Ø., 2010. A protocol for generating a high-quality genome-scale metabolic reconstruction. Nature Protocols 5, 93–121.

Thorleifsson, S. G., Thiele, I., 2011. rBioNet: A COBRA toolbox extension for reconstructing high-quality biochemical networks. Bioinformatics 27 (14), 2009–2010.

Vlassis, N., Pacheco, M. P., Sauter, T., Jan. 2014. Fast reconstruction of compact context-specific metabolic network models. PLoS Comput Biol 10 (1), e1003424.

von Kamp, A., Klamt, S., Jan. 2014. Enumeration of smallest intervention strategies in genome-scale metabolic networks. PLoS Comput Biol 10 (1), e1003378.

Wilhelmy, L., 1850. Ueber das Gesetz, nach welchem die Einwirkung der Säuren auf den Rohrzucker stattfindet (The law by which the action of acids on cane sugar occurs). Poggendorff's Annalen der Physik und Chemie 81, 413–433.

5. Supplementary Material

5.1. Numerical Linear algebra

5.1.1. Rank Computation

To compute the rank of a matrix, we employed the sparse LU factorisation package LUSOL (Gill et al., 1987, 2005). Given a sparse matrix $A \in R^{m \times n}$, LUSOL computes factorisations of the form

$$P_1 A P_2 = LDU, \tag{10}$$

where P_1 and P_2 are permutations, L is lower trapezoidal with unit diagonals, D is diagonal and nonsingular, U is upper trapezoidal with unit diagonals, and the rank of each factor L, D, U is $r \leq \min(m, n)$. This is LUSOL's estimate of rank(A).

The permutations are chosen to keep L and U sparse, subject to bounds on the off-diagonal elements of L and U. Threshold Partial Pivoting (Gill et al., 1987) requires $|L_{ij}| \leq \tau$ for some tolerance $\tau \in (1, 10]$, where $\tau = 2$ is a reasonable choice. Threshold Rook Pivoting (Gill et al., 1987) also requires $|U_{ij}| \leq \tau$ and is likely to be more reliable on general sparse A. We have observed that net stoichiometric matrices $S = R - F$ have a sharply defined rank that can be estimated reliably by Threshold Partial Pivoting on either S or S^T and this is cheaper than applying Threshold Rook Pivoting. The same is true for estimating the rank of $[F, R]$.

5.1.2. Identification of dependencies

To investigate the rationale for row rank-deficiency of $A := [F, R]$ derived directly from a reconstruction, we used Threshold Partial Pivoting to obtain the factorisation (10). The row permutation P_1 partitions the rows as

$$P_1 A = \begin{bmatrix} B \\ C \end{bmatrix},$$

where $B \in R^{r \times n}$, $C \in R^{(m-r) \times n}$, and rank(B) = r. By definition of rank, the over-determined linear system $B^T W = C^T$ is consistent (has a solution W). The nonzero entries of each column of W reveal the dependencies between rows of B and C. We obtained W by solving $\min \|B^T W - C^T\|$ using sparse QR factorisation (`W = B'\C';` in MATLAB). The column permutation P_2 further partitions A as

$$P_1 A P_2 = \begin{bmatrix} B_1 & B_2 \\ C_1 & C_2 \end{bmatrix},$$

where B_1 is $r \times r$. We could obtain W more efficiently by solving $B_1^T W = C_1^T$, where $B_1 = L_1 U_1$ is already factorised in terms of the first r rows and columns of L and U. The nonzero entries of W reveal the dependencies between rows of B_1 and C_1. As each row of B_1 and C_1 corresponds to a different molecular species, one can use W to investigate the biochemical rationale for dependency among rows of $[F, R]$ if dependency is observed.

5.2. Combinatorial dependence in exceptional models derived from genome-scale reconstructions

Combinatorial dependence among the rows of [F, R] implies that [F, R] is row rank deficient. However the reverse implication is not necessarily true. This is because linear dependence depends on the nonzero numerical values of elements in [F, R] whereas combinatorial dependence depends only on the sparsity pattern of [F, R] and not the numerical values in [F, R]. Nevertheless, it is of interest to check if an [F, R] that is row rank deficient also contains combinatorially dependent rows.

Only 3 of the 29 reconstructions subjected to the four conditions in Section 2.4.4 resulted in a row rank-deficient [F, R], with rank at most 3 lower than the number of rows. If rank([F, R]) is less than the number of rows of [F, R] then numerical linear algebra (cf. Section 5.1.2) can be used to test for dependency between rows to identify possible reasons for the row rank deficiency. The three models with rank-deficient [F, R] were from compartmentalised genome-scale models. In each of the three models, one could find at least one dependency between a dependent molecular species (one dependent row of [F, R]) and a set of independent molecular species within the same sub-cellular compartment (a set of linearly independent rows of [F, R]).

Each dependency in [F, R] was due to the existence of two disjoint sets of molecular species, one having a cofactor moiety in common and one having a non-cofactor moiety in common, with one cofactor and one non-cofactor always present in reactant complexes within that sub-cellular compartment. That is, neither the cofactor moiety nor the non-cofactor moiety was synthesised or degraded within that sub-cellular compartment. This reflected one or more reactions that were missing from the original reconstruction that would either synthesise or degrade the moiety within that sub-cellular compartment. Such reactions would give [F, R] full row rank, as inevitably there would be at least one reaction where the cofactor and non-cofactor moiety would not both be represented within a reaction complex at the same time. Table 1 illustrates a specific example of such a dependency within a model derived from a *Saccharomyces cerevisiae* reconstruction (iMM904). Alternatively, the dependency reflected the omission of a reaction to transport either the cofactor or non-cofactor moiety into that sub-cellular compartment. Such a reaction would typically not simultaneously involve both the cofactor moiety and the non-cofactor moiety rendering [F, R] of full row rank.

Table 1: An example of a set of endoplasmic reticulum reactions within a genome-scale *Saccharomyces cerevisiae* reconstruction (iMM904), that results in a rank deficient $[F, R]$ even when all molecular species are stoichiometrically consistent and all reactions are net flux consistent (given exchange reactions, assuming each reaction is reversible and omitting nontrivial rows). Besides the 6 reactions illustrated, Phytosphingosine, Sphinganine, Tetracosanoyl-CoA, Hexacosanoyl-CoA and Phosphate are not involved in any other reactions within the endoplasmic reticulum. Phytosphingosine and Sphinganine form one set reactions, while Tetracosanoyl-CoA, Hexacosanoyl-CoA and Phosphate form another set. These sets are disjoint as they do not share a molecular species in common. Observe that the support of both disjoint sets is identical, i.e., all reactions contain at least one member of both sets in the same reaction complex. This renders the corresponding rows of $[F, R]$ row rank deficient. In fact, the rank of these 5 rows is 4, hence leading to row rank deficiency of $[F, R]$. A more comprehensive model would have the 3-ketodihydrosphingosine reductase reaction (Phytosphingosine + NADPH $\rightleftharpoons$ 3-Dehydrosphinganine + NADP) that would result a full row rank $[F, R]$.

R1 alkaline ceramidase (ceramide-1)
R2 alkaline ceramidase (ceramide-1)
R3 alkaline ceramidase (ceramide-2)
R4 alkaline ceramidase (ceramide-2)
R5 sphingoid base-phosphate phosphatase (sphinganine 1-phosphatase)
R6 sphingoid base-phosphate phosphatase (phytosphingosine 1-phosphate)
A1 CERASE124er
A2 CERASE126er
A3 CERASE224er
A4 CERASE226er
A5 SBPP1er
A6 SBPP1er

	Reaction Name	R1	R2	R3	R4	R5	R6
Metabolite species name	Abbreviation	A1	A2	A3	A4	A5	A6
Ceramide-1 (Sphinganine:n-C24:0)	cer1_24[r]	-1	0	0	0	0	0
Ceramide-1 (Sphinganine:n-C26:0)	cer1_26[r]	0	-1	0	0	0	0
Ceramide-2 (Phytosphingosine:n-C24:0)	cer2_24[r]	0	0	-1	0	0	0
Ceramide-2 (Phytosphingosine:n-C26:0)	cer2_26[r]	0	0	0	-1	0	0
Coenzyme A	coa[r]	-1	-1	-1	-1	0	0
Proton	h[r]	-1	-1	-1	-1	0	0
Water	h20[r]	0	0	0	0	-1	-1
Sphinganine 1-phosphate	sph1p[r]	0	0	0	0	-1	0
Phytosphingosine 1-phosphate	psph1p[r]	0	0	0	0	0	-1
Phytosphingosine	psphings[r]	0	0	**1**	**1**	0	**1**
Sphinganine	sphgn[r]	**1**	**1**	0	0	**1**	0
Tetracosanoyl-CoA	ttccoa[r]	**1**	0	**1**	0	0	0
Hexacosanoyl-CoA (n-C26:0CoA)	hexccoa[r]	0	**1**	0	**1**	0	0
Phosphate	pi[r]	0	0	0	0	**1**	**1**

5.3. Reproduction of numerical results

All of the reconstructions and code required to reproduce the numerical results referred to in this paper are publicly available within the COBRA toolbox (Schellenberger et al., 2011) via `https://github.com/opencobra/cobratoolbox`. The steps are as follows:

1. Install MATLAB version 8.4.0.150421 (R2014b) or above. Earlier versions of MATLAB may also suffice, but have not been tested for this purpose.
2. Install the latest version of The COBRA toolbox (more recent than December 1, 2015). From a unix command line, enter the command:
   git clone https://github.com/opencobra/cobratoolbox.git
3. Optionally install a 64-bit Unix implementation of LUSOL
   `http://stanford.edu/group/SOL/software/lusol/`.
   From a Unix command line, enter the command
   git clone `https://github.com/nwh/lusol.git`.
   Installation is optional as otherwise the sparse LU factorization provided in MATLAB is employed.
4. The folder `cobratoolbox/testing/testModels/modelCollectionFR` contains each of the reconstructions in COBRA Toolbox format (one MATLAB .mat file for each reconstruction). Each .mat file was derived from the original SBML file that was published with the respective papers or provided as published updates to the original SBML file, as detailed within the function
   `cobratoolbox/testing/testModels/modelCitations.m`.
5. All numerical results can be reproduced by calling the MATLAB function
   `cobratoolbox/papers/Fleming/FR_2015/checkRankFRdriver.m`.
   This driver file passes each reconstruction to `checkRankFR.m`, which generates the corresponding model as detailed in Section 2.4.4 and uses numerical linear algebra to check rank([F, R]), as described in Section 5.1.1.

Supplementary references

References

Gill, P. E., Murray, W., Saunders, M. A., Wright, M. H., Apr. 1987. Maintaining LU factors of a general sparse matrix. Linear Algebra and its Applications 88–89, 239–270.

Schellenberger, J., Que, R., Fleming, R. M. T., Thiele, I., Orth, J. D., Feist, A. M., Zielinski, D. C., Bordbar, A., Lewis, N. E., Rahmanian, S., Kang, J., Hyduk, D., Palsson, B. Ø., 2011. Quantitative prediction of cellular metabolism with constraint-based models: the COBRA Toolbox v2.0. Nature Protocols 6 (9), 1290–1307.